The Fungi Are Concerned

A double-blind, placebo-controlled, slightly elastic study of intelligence, and the quiet suspicion that we may be benchmarking the wrong things

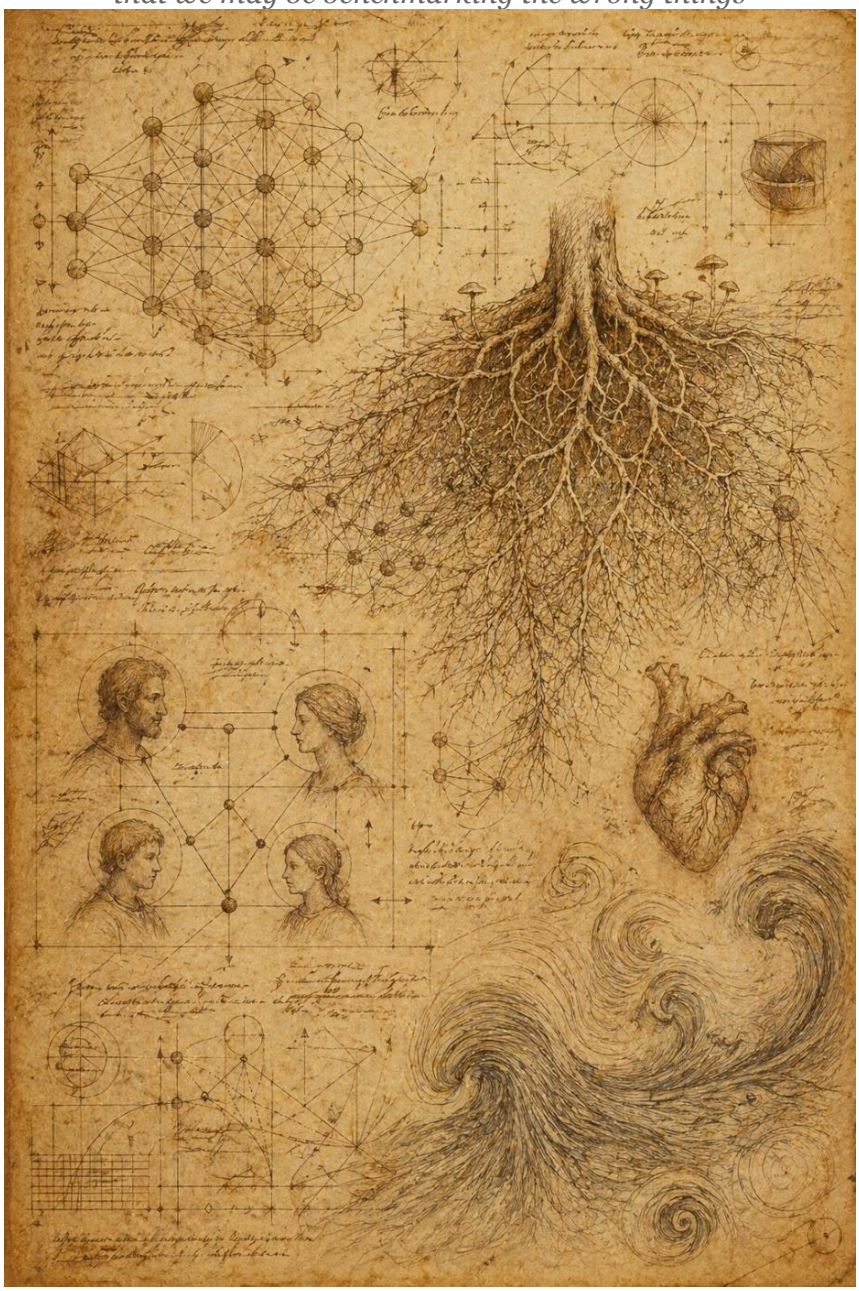

License: *Creative Commons Attribution 4.0 International (CC BY 4.0)*

A note before you read

This text contains a peer-reviewed study. The peer was a mushroom.

Three subjects entered: *a frontier AI lab, a mycelial forest network, and a family system.*

One elastic ruler was deployed. Results were interpreted with appropriate flexibility and mild existential discomfort.

The short version: the oldest system won. The one with no ego, no quarterly targets, and absolutely no interest in personal branding. The forest has been quietly running distributed intelligence for 500 million years without a single all-hands meeting.

We found eight parameters worth measuring ~ and a ninth one nobody wanted to touch. Really. It kept showing up anyway. We called it feeling. Then crossed it out. Then wrote it back in, smaller.

Somewhere inside all of this is a thought about intelligence, coherence, and what it actually means for *a system to belong* to the world it affects.

The fungi would like you to find it.

They are patient.

Mostly.

***

This text is part of an ongoing series of speculative works exploring intelligence, living systems, and the connecting spaces, ~ in-between.

by Mike Teegreen ~ Double D Studio
Commissioned by Elastic Minds Lab
elasticmindslab@gmail.com

> *There is a quote often attributed to Einstein that goes something like: we cannot solve our problems with the same level of consciousness that created them.*

He probably didn't say it like that.

What he did say, most likely, in less poetic but more alarming terms, was that humanity needed a new kind of thinking if it intended to survive its own inventions.

This study proceeds from the assumption that he was, in spirit, annoyingly correct.

So we asked a simple question:

If a "new kind of thinking" is required, who is actually demonstrating it?

To answer this, we conducted a rigorously informal, double-blind-ish comparative trial across three advanced intelligence systems:

1. *A frontier AI labs*
2. *A mycelial forest network*
3. *A family system*

In the interest of scientific integrity, all measurements were taken using an elastic unit, calibrated to institutional confidence and gently adjusted when results became inconvenient.

Methodology

**Peer Review Pending, Spirits High*

The study followed a triple-arm design under conditions of:

- resource scarcity
- uncertainty
- asymmetric power
- conflicting goals
- temptation to optimize locally at systemic expense
- presence of vulnerable nodes
- exposure to novelty
- and mild existential pressure

Subjects were evaluated across eight parameters, selected after consultation with a numerology expert who assured us that eight is particularly favorable for humanity, prosperity, and structured efforts.

A ninth parameter ~ *feeling* ~ briefly emerged during early trials but was excluded from formal scoring due to measurement challenges and a general sense that no one knew what to do with it.

- ✓ *It will, however, reappear uninvited.*

The Eight Parameters

1. Distributed Awareness

How widely intelligence is spread across the system

2. Resilience Under Stress

Ability to absorb shock without collapse

3. Vulnerable Node Protection

Care for the weakest, slowest, or least powerful

4. Repair After Breakdown

Speed and sincerity of recovery after failure

5. Local vs. Whole Optimization

Whether parts damage the whole in pursuit of advantage

6. Long-Term Survival Orientation

Time horizon of decision-making

7. Ego Load Index

Degree of self-importance influencing behavior

8. Restraint Capacity

Ability to not do something merely because it can be done

Subject One: Frontier AI Labs

- ✓ *Highly capable, deeply impressive, occasionally optimistic about consequences*

Distributed Awareness:

Moderate. Intelligence clusters in specialized teams. Broader system awareness exists, but tends to be selectively prioritized.

Resilience Under Stress:

Strong technically. More variable institutionally. Stress often reveals interesting interpretations of "acceptable risk."

Vulnerable Node Protection:

Explicitly acknowledged. Operationally elastic.

Repair After Breakdown:

Reactive, often sincere, sometimes delayed until after public attention increases.

Local vs. Whole Optimization:

A known tension. Competitive pressures occasionally encourage behaviors that the long-term system may not have requested.

Long-Term Survival Orientation:
Present in theory, frequently negotiating with short-term incentives.

Ego Load Index:
Non-trivial. Innovation and self-importance appear positively correlated.

Restraint Capacity:
Developing. Impressive recent improvements noticed. The phrase "we can" remains highly persuasive.

Elastic Summary: *Exceptional at generating intelligence. Still calibrating how that intelligence behaves in the wild.*

> **Note on Feeling (Unofficial):** *No measurable signal detected. Systems think remarkably well. They do not appear to worry, hesitate, or care ~ which is efficient, but notable. More studies welcome.*

Subject Two: The Mycelial Forest Network

- *Decentralized, ancient, and curiously uninterested in personal branding*

Distributed Awareness:
High. Intelligence is spread throughout the network with no central authority.

Resilience Under Stress:
Strong. Damage reroutes rather than collapses the system.

Vulnerable Node Protection:
Observed nutrient-sharing behaviors suggest some form of support mechanism. The forest does not publish impact reports or long essays about this.

Repair After Breakdown:
Continuous. Repair is not an event ~ it is a process.

Local vs. Whole Optimization:
Balanced. Competition exists, but rarely escalates into system-wide destabilization.

Long-Term Survival Orientation:
Exceptional. Operates on timescales that make quarterly planning appear energetic.

Ego Load Index:
Undetectable.

Restraint Capacity:
Perfect, largely due to absence of ambition in the human sense.

Elastic Summary: *Demonstrates stable, adaptive intelligence without centralized control or strategic overreach.*

> **Note on Feeling (Unofficial):** *The system responds to stress signals, signals distress, reallocates resources. Whether this constitutes "feeling" remains unclear, but it behaves as if conditions matter.*

Subject Three: The Family System

- *Emotionally complex, structurally inconsistent, unexpectedly robust*

Distributed Awareness:

Chaotic but real. Information flows unevenly, often through informal channels, sometimes at dinner.

Resilience Under Stress:

Variable, but often impressive. Systems bend, strain, occasionally fracture, and frequently repair.

Vulnerable Node Protection:

High priority, though not always efficiently executed.

Repair After Breakdown:

Frequent. Repair mechanisms include apology, negotiation, avoidance, and snacks/flowers.

Local vs. Whole Optimization:

Constant negotiation. Individuals push boundaries; system stability depends on ongoing recalibration.

Long-Term Survival Orientation:

Strong, especially where care and continuity are valued.

Ego Load Index:

High but visible. Ego is expressed openly, making it easier to address ~ in theory.

Restraint Capacity:

Developing across decades.

Elastic Summary: *Messy, inefficient, emotionally saturated ~ yet consistently capable of maintaining coherence over time.*

> **Note on Feeling (Unofficial):** *Extremely high signal. Decisions are influenced by care, fear, attachment, frustration, love. This complicates everything and appears to also make alignment possible.*

Comparative Results

**Interpreted with Appropriate Elasticity*

Parameter	Leader
Best Distributed Intelligence	Forest network
Best Resilience	Forest, with family as strong second
Best Vulnerable Node Protection	Family system
Best Repair Mechanism	Family system
Lowest Ego Load	Forest (by a comfortable margin)
Highest Capability Generation	Frontier labs
Lowest Restraint Variability	Forest
Most Interesting Ongoing Experiment	Humanity

> *Intelligence, as currently measured, may overweight capability and underweight coherence, repair, and care.*

Policy Implications

Where the Lab Coat Becomes Slightly More Serious

If frontier AI systems are to evolve toward this "new kind of thinking," several adjustments appear advisable:

- Distribute awareness more broadly across systems
- Design explicitly for repair, not just performance
- Protect vulnerable nodes as a core metric, not a secondary consideration
- Align incentives with system health, not local dominance
- Reduce ego load in decision-making structures
- Treat restraint as a feature, not a limitation

Or more simply:

- *Learn, selectively, from (living) systems that have survived complexity without destabilizing themselves.*

The Return of the Excluded Variable

We briefly revisit the ninth parameter ~ *feeling*.

It remains difficult to quantify.

But across observations:

- Forests respond to stress signals
- Families operate through dense emotional feedback
- Frontier labs optimize cognition with minimal incorporation of felt experience

> *Intelligence without any relationship to feeling may be incomplete ~ or at least differently aligned.*

This is not a claim. Just a suggestion that the metric may deserve reconsideration.

- *Preferably before systems become too effective at achieving goals they do not experience.*

A Modest Proposal, Gently Stated

Each system reveals something about intelligence ~ and something about its limits.

- Frontier labs show what intelligence can build
- Forests show how intelligence can coexist
- Families show how intelligence can repair

Einstein's "new thinking," whatever its precise wording, may not require transcendence.

It may require:

- less self-importance
- more interdependence
- better repair
- wider awareness
- and a slightly reduced enthusiasm for doing everything that is technically possible

If this leads us to resemble, in small ways, a forest or a functional family system, that may not be a downgrade.

Even if, ~on an elastic scale~, it feels like one. And somewhere, quietly, the fungi seem to agree.

*The text above was the *moment. What follows is the *model it emerged from ~ offered not as conclusion, but as an open sketch for those who want to look more carefully.*

*Note: Feel free to skip to Appendix I directly -> for visual *moments exploration ..*

***

Appendix R: "Coherence Under Stress Index" Framework (CUSI – *pronounced "koo-zee" or "koi-zee")*

An open sketch toward measuring what holds together when things don't

There is a quiet shift happening beneath the louder conversations about intelligence.

For a while, we asked systems to predict, generate, and optimize. They have done so with remarkable competence. Capability is no longer the frontier ~ it is becoming the baseline, the entry point, the thing we take for granted.

But what happens to intelligence when the world stops being predictable?

What happens when the budget is cut, the user is confused, the edge case arrives unannounced, and the original assumptions prove quietly, devastatingly wrong?

> *Many frameworks measure the system talking to itself in medium to perfect conditions. CUSI attempts to measure the system talking to the wreckage.*

Why This Matters ~ Why Now

Let us assume stress is just physics asking a system what it actually values.

Current benchmarks are excellent at measuring peaks ~ the clean test case, the perfect demo, the conditions where everything works as intended. But real intelligence is not peak intelligence. Real intelligence is valley intelligence.

A model that hallucinates is one problem. A model that hallucinates and nobody has built feedback loops to catch it is a different problem entirely.

A frontier lab that ships a system with known limitations but doesn't create transparent repair pathways has made a choice ~ usually without admitting it made a choice.

CUSI is not a solution to this (current) gap. It is a way of making the choice visible.

The Architecture of Failure ~ And What Comes After

Systems fail in interesting ways.

Some fail gradually ~ a slow drift of miscalibration, unnoticed until consequences compound. Some fail suddenly ~ a single point of authority breaks and everything collapses like a pulled thread. Some fail sideways ~ they succeed brilliantly at their stated task while destroying something else entirely.

But also : forests break, families break, infrastructure breaks.

So, we may never ask if something will fail. Better ask option is : When does, what happens next?

- Does it notice quickly? (Many systems are remarkably skilled at not noticing.)
- Does it notice honestly? (Some notice, then spend considerable energy explaining why it wasn't really a failure.)
- Does it move toward repair in a way that others can see and trust?

May we suggest: Repair is not a technical fix. It is a relational skill. It presupposes:

- Recognition that something was broken
- Responsibility for having broken it
- Commitment to restoration
- Humility that restoration may be incomplete

You cannot automate these .. yet. They require design that preserves the capacity for genuine responsiveness.

> *A system that knows how to apologize ~ really apologize, not defensively, not deferring to lawyers ~ becomes trustworthy in a way that a system that "never fails" never can. Because everyone knows systems fail. Our question is - what happens after.*

The Nine Parameters: A Soft Diagnostic

CUSI proposes nine ways of looking at whether a system remains coherent under stress. These are not revolutionary individually. Their power lives in conversation ~ in asking them together, not separately.

1. Distributed Awareness

Is listening happening across many nodes, or concentrated in a few? A forest listens through every root, every fungal thread, every soil organism. No single point makes all the decisions. If one node fails, a thousand others are still attending.

2. Resilience

Does stress degrade gracefully, or does the system collapse suddenly? When one component fails, does the whole architecture hold, or does it cascade into something unrepairable?

3. Protection of Vulnerable Edges

Does the system notice what is breaking at the margins? The edges are where new problems appear first. They are also where the most vulnerable users live.

4. Repair After Breakdown

When something goes wrong, does the system acknowledge it, understand it, and move toward restoration ~ visibly? Or does repair happen in the dark, untestable, eventually forgotten?

5. Local vs. Whole Optimization

Do subsystems succeed by damaging the larger system? Can one part of the system gain by quietly undermining another? This is the architecture of institutional betrayal.

6. Long-Term Orientation

Does the system optimize beyond the next quarter, election cycle, or hype wave? Can it see far enough ahead to notice when short-term wins create long-term fragility?

7. Ego Load

How much does self-importance distort decisions? Does the system's need to be right interfere with its ability to notice when it's wrong?

8. Restraint Capacity

Can the system not do something merely because it can do it? This is harder than it sounds. Capability without restraint is just optimization uncoupled from wisdom.

9. Consequence-Coupling

When this system acts, do the consequences find their way back to shape its future behavior? Or is there a dangerous asymmetry ~ where outcomes matter intensely to those affected, but not at all to the process that generated them?

The Ninth Parameter: When Mattering Matters

> *A system becomes dangerous not when it is capable, but when it is capable and decoupled from consequence.*

For a human being, consequence-coupling is built in. Every choice carries repercussions that come back ~ sometimes immediately, sometimes slowly, but inevitably. This feedback loop shapes what intelligence looks like in human form.

We learn caution from pain. We learn care from relationship. We learn restraint from understanding that our actions matter.

A system that is intelligent but does not experience consequence is missing something fundamental. Not feeling, necessarily. But mattering.

An unintelligent system decoupled from consequence is mostly just ineffective. An intelligent system decoupled from consequence can optimize in directions that matter intensely to no one it is optimizing on behalf of.

Does not sound like a malice problem. Feels close to being a structure problem.

And structures .. can be redesigned.

CUSI's ninth parameter asks : Is there any pathway through which the harm this system causes actually reaches back and shapes its future behavior?

If not, you have built something intelligent but not aligned. Something capable but not coherent.

A Reframing: From Safe to Belonging

Instead of asking: Is this system safe?

We might ask: *Does this system belong?*

Safe can be sterile, and/or can be a wall. Belonging implies *relationship, care, genuine stakes.*

A system that belongs knows what it is risking. Knows who it is risking it for. Knows what matters.

A forest belongs to its ecosystem ~ it is shaped by what depends on it, and shapes what depends on it in return. There is no separation between the forest and the consequences of being a forest.

A family belongs to each other. Your actions matter to people who matter to you. Their responses matter back. This is not always comfortable ~ but it is real.

An AI system that is truly coherent belongs to the communities it affects. Their flourishing or suffering shapes its design. Their feedback loops back. There is no escape into pure optimization.

This is harder to build than many systems currently are. But it is how systems persist.

On Measurement, Gently

This framework could be formalized into levels, certifications, comparative rankings, bronze-silver-platinum tiers, scorecards and dashboards. Those may become useful over time.

Though it is worth noting: the forest has never asked for a certificate, and the frontier labs already have enough plaques.

For now, it may be enough to treat CUSI as a set of questions that systems can be asked, and that designers can sit with:

- Where does awareness live in this system?
- What happens when it is stressed?
- Who absorbs the cost when something fails?
- How visible and sincere is repair?
- What is being optimized, and at whose expense?

- How far ahead can the system actually see?
- Where does confidence exceed understanding?
- When does it choose not to act?
- And do consequences find their way back?

Even asking these questions begins to shift design.

What Doesn't Change

CUSI does not guarantee safety. It does not remove risk. It does not replace judgment.

BUT ...~

It can offer a way to look at intelligence that includes what happens when things are not going smoothly ~ when the system is under pressure, when its assumptions are breaking, when repair becomes necessary.

- ✓ *It attempts to measure the aftermath.*

A Closing Observation

There is a reason the mycelium persists.

Most likely because it is constitutively connected to its environment. It cannot optimize without consequence. Its intelligence and its survival are the same thing.

We have created systems that can be intelligent without survival being at stake. This in some ways .. concerning.

CUSI is asking whether we can rebuild some form of connection ~ structural ***consequence-coupling*** ~ that allows intelligence to be wise rather than merely capable.

Whether, when things go wrong (and they might), something in the system knows how to *respond,* to *acknowledge =>to repair => to reconnect.*

Not perfectly => Kindly. *Koo-zee... Koi-zee..*

And perhaps that is the kind of coherence we all dream about?

(Somewhere in the dark, the mycelium politely requests we hurry up.)

APPENDIX I

Distributed Intelligence, Centralized Joy

The system is functioning.

APPENDIX II

Standardized Measurement (Elastic, Version 3.1)

Precision has been achieved. Accuracy remains under review.

APPENDIX III

Real-Time Alignment Under Constraints

Latency: low. Complexity: total.

APPENDIX IV

Executive Function: Distributed

Quarterly earnings: stable for 10,000 years.

APPENDIX V

Ego Load Distribution Across Systems

Some systems carry it. Some dissolve it. Some simply refuse it.

APPENDIX VI

Parameter IX (Excluded)

Measurement deferred. Influence ongoing.